arXiv:1704.08417v1 [math.OC] 27 Apr 2017

EFFICIENT PROJECTION PARTITIONING FOR PARALLEL MULTI-OBJECTIVE INTEGER OPTIMISATION

WILLIAM PETTERSSON AND MELIH OZLEN

Abstract. This paper introduces a new method of partitioning the solution space of a multi-objective optimisation problem for parallel processing, called Efficient Projection Partitioning. This method projects solutions down into a single dimension, greatly reducing the cost of partitioning the search space. We test EPP on a variety of randomly generated multi-objective combinatorial optimisation problems. The results are compared with the state of the art in parallel partitioning, and we show that in all scenarios tested, our new algorithm performs significantly better. Our proposed method allows the generation of non-dominated sets of larger problems with more decision variables or objective functions through the use of highly parallel computational infrastructure. Source code is provided to allow others to utilise, build upon and improve the algorithm.

1. Introduction

In multi-objective integer programming (MOIP), one must consider a range of objective functions. One method of exactly solving a MOIP problem is computing all the non-dominated solutions. This set of non-dominated solutions can be called the Pareto set, or the Pareto frontier. A decision maker can use such a set to compare the various trade-offs that can be made between the objective functions.

Finding the Pareto set to a MOIP is computationally expensive, and many multi-objective combinatorial optimisation problems have been shown to be NP hard. To alleviate this, some algorithms will use (meta-)heuristics, or swarming [1, 2] or evolutionary [3, 4] algorithms to compute a set which approximates the Pareto set. However, this paper will only consider algorithms which calculate the exact Pareto set, with no omissions or inaccuracies. For a complete survey on multi-objective optimisation in general, see [5], and for a very recent and thorough look at exact MOIP algorithms, focusing on branch and bound algorithms, see [6].

The solutions in such a Pareto set are distributed in the solution space, so recent advances in the parallelisation of exact MOIP algorithms has involved partitioning this space. This can be thought of as a generalisation of the bi-objective two-phase method (see [5] for details). The bi-objective two-phase method first finds a set of supported solutions, and then uses this set of solutions to split the search space into smaller partitions which can be searched in parallel. This process of first determining all partitions and then solving all subproblems is seen in [7, 8]. Other algorithms use a more recursive approach, where more subproblems are created as new solutions are found [9, 10, 11]. Many of these algorithms suffer from a common (and well-known) drawback, however. They must calculate the nadir vector of the solution set, an expensive calculation for problems with many objectives, to partition the solution space.

This paper describes a new method of partitioning the solution space which we call the Efficient Projection Partitioning (EPP) method. Solutions are projected into a single dimension, where upper and lower bounds can be calculated much faster. This avoids the expensive computation of the nadir vector, and the simplicity of the design makes implementation straight-forward and less error-prone. A wide variety of tests show that EPP outperforms existing state-of-the-art algorithms, and also scales far better as more objective functions are introduced. These improvements significantly reduce the cost of solving difficult multi-objective combinatorial optimisation problems.

1.1. Paper organisation.

A brief background to multi-objective integer optimisation in general, as well as parallel computing in multi-objective optimisation, is given in Section 2. Section 3 gives a brief overview of some existing algorithms in the literature, before introducing our new algorithm in Section 3.2. We give implementation and testing details in Section 4. Test results and a discussion follow in Section 5, along with a conclusion in Section 6.

2010 *Mathematics Subject Classification.* Primary: 90C29; Secondary: 90C10, 68W10.

Key words and phrases. multi-objective optimisation, parallel computing, integer programming, combinatorial optimisation.

This study was supported by the Australian Research Council (DP140104246).

2. Background

2.1. **Multi objective optimisation.** We give a brief introduction to multi-objective integer optimisation here. For a more complete background on multi-objective integer optimisation algorithms, see e.g. [5, 6]. For conciseness we assume that all objectives are to be minimised.

Let $X = \{x | x \in \mathbb{Z}^n \text{ and } Ax \leq b\}$ be the feasible space for a MOIP problem, where A and b are integer-valued matrices of appropriate size, and let $\{f_1, \ldots, f_K\}$ denote the K objective functions for said problem. We assume that these objective functions have integer coefficients.

Definition 1. Given two distinct solutions y, z to a MOIP problem, we say that y dominates z if

(1) $f_k(y) \leq f_k(z)$ for all $k \in \{1, \ldots, K\}$, and
(2) $f_k(y) < f_k(z)$ for at least one $k \in \{1, \ldots, K\}$.

For a subset $S \subseteq \{1, \ldots, K\}$, we will write $f_S(y) < f_S(z)$ if y dominates x when the objective functions indicated by S are considered. For brevity, we will also write $f_{[K]}(y) < f_{[K]}(z)$ for an integer K where the whole set $\{1, \ldots, K\}$ is considered.

We follow the notation of [7] and define Opt^K to be the solution set to a MOIP problem with K objective functions.

Definition 2. $\text{Opt}^K = \{z | z \in X \text{ and } \nexists\, y \in X \text{ s.t. } f_{[K]}(y) < f_{[K]}(z)\}$

This set of solutions is called *non-dominated* (or sometimes efficient, Pareto optimal, or simply Pareto).

One way of creating sub-problems of a given MOIP problem is to only consider some of the objective functions. The set Opt_S^l, where S is a subset of $\{1, \ldots, K\}$ of size l, is the set of solutions to the problem when only the l objectives indicated by S are considered.

Definition 3. $\text{Opt}_S^l = \{z | z \in X \text{ and } \nexists\, y \in X \text{ s.t. } f_S(y) < f_S(z)\}$

Partitioning algorithms tend to require the calculation the extremal points of the search space, called the *ideal* and *nadir* vectors. These are vectors in the solution space which respectively take on the lowest and highest values of each individual objective function across all Pareto solutions.

Definition 4. The ideal vector of a MOIP problem is given by $I = (i_1, \ldots, i_K)$ where $i_j = \min\{f_j(y) | y \in \text{Opt}^K\}$ for all $j \in \{1, \ldots, K\}$.

Definition 5. The nadir vector of a MOIP problem is given by $I = (i_1, \ldots, i_K)$ where $i_j = \max\{f_j(y) | y \in \text{Opt}^K\}$ for all $j \in \{1, \ldots, K\}$.

2.2. **Parallel computing.** Parallel computing allows multiple calculations to be completed in parallel (i.e. at the same time). We will use the term *computing thread*, or simply *thread*, to represent one logical unit that is able to perform calculations independently. The same term is also used in software. In hardware, one logical unit is often called a core. Modern phones are often advertised with two or four cores, and desktops can have eight or more cores. Supercomputing facilities often far exceed these numbers, and recent advances in CPU architecture has realised processors with 72 or more cores.

2.3. **Objective space splitting and MOIP solvers.** Parallel algorithms which split up the search space rarely specify the exact solver to use. Instead, the algorithms will create a number of sub-problems and solve each independently. This means that any comparison of algorithms must also consider the choice of solver. Many papers utilise a solver specialised for their given combinatorial optimisation problem when testing a partitioning algorithm. This is ideal when trying to solve a specific problem, but does present problems when comparing algorithms. To alleviate this, we will use the same generic MOIP solver for all testing. This does mean that our results will most likely not line up exactly with the results obtained by others. However our aim is to compare and analyse the differences in a generic and broader MOIP setting, not just for a specific combinatorial optimisation problem.

Note that the sub-problems are generated from the partitioning of the search space, and thus solving such sub-problems may require adding upper and/or lower bounds on the values of the objective functions. Some solvers do not allow the addition of lower bounds, which could require an adjustment to the partitions created (see e.g. Section 3.4.2 in [7]). Even if the addition of bounds is possible, there is also always the possibility that these additional constraints will slow down the algorithms significantly. Again, this will all depend on the choice of solver, but these are some of the factors to consider when selecting or designing a solver for use in a parallel partitioning algorithm.

We use MOIP-AIRA [12, 13] as our MOIP solver. It is a state of the art and open source exact solver which is not optimised for any one specific problem, and therefore can easily and quickly be integrated into the partitioning algorithms we test. The design of MOIP-AIRA makes the addition of upper bounds

essentially a zero-cost operation. Lower bounds can also be added without requiring the addition of extra constraints to the problem or bounds on any decision variables. Exactly how this is achieved is explained in Section 3.1. However this does mean that, given some set of bounds, MOIP-AIRA will find all non-dominated solutions within these bounds but also potentially some solutions outside these bounds. This could be considered a cost, or penalty, of using MOIP-AIRA, however we believe this to still be most efficient as experiments have shown that adding actual lower bounds adds significantly more to the running time than finding the extra solutions.

3. Algorithms

3.1. **MOIP-AIRA.** MOIP-AIRA is the chosen MOIP solver for our experimentation. We briefly detail our modifications which constrained it to a given partition. For a thorough and formal description, see [12].

MOIP-AIRA begins with upper bounds on all objectives initially set to ∞. Then, MOIP-AIRA solves the lexicographic version of the problem, constrained to these upper bounds, to find solutions. As it finds solutions, the upper bounds are slowly reduced until the problem is infeasible.

By starting these upper bounds at the boundary of the partition instead of infinity, we have added an upper bound to the sub-problem at zero cost. Testing showed that adding lower bounds as actual constraints to the sub-problem was very inefficient. Instead, MOIP-AIRA is run as designed, but every time the upper bounds are changed, they are also compared to the lower bounds. When all upper bounds are reduced below the lower bounds, all relevant solutions within a given partition have been found and the algorithm will terminate.

3.2. **Efficient Projection Partitioning.** The solution space for a MOIP can be envisioned as a k-dimensional vector space, where each dimension represents one objective function. The Efficient Projection Partitioning (EPP) algorithm projects the whole solution space down to one dimension. Given a solution vector $\mathbf{x} = (x_1, \ldots, x_k)$, the projection is achieved through the k-th projection map $proj_k(\mathbf{x}) = x_k$. That is, the solution space is partitioned by only considering the values attained by one objective function. This allows EPP to partition the space without first calculating the complete nadir vector, giving EPP an advantage over many other partitioning algorithms. First we need the following lemma.

Lemma 1. *Any solution in* Opt^K *which achieves a maximum value on* f_k *is also a solution in* $Opt^{K-1}_{\{1,\ldots,k-1\}}$.

For a proof of this Lemma, see e.g. Lemma 4.1 in [14] or Theorem 2 in [7]. EPP calculates $\mathrm{Opt}^{K-1}_{\{1,\ldots,k-1\}}$ recursively, and then uses the above lemma to compute the maximum value of f_k. A lower bound on f_k is trivially found using a standard single-objective IP solver. This gives a range of values which f_k can take, which is divided up equally amongst all threads. Figure 1 provides a visual representation of a sample partition for a 3-objective problem.

Data: A K-objective MOIP, and an integer T representing number of threads to use.
Result: The non-dominated solutions.
if $K = 1$ **then**
  Solve the single-objective problem and return Opt^1_1
else
  Calculate $\mathrm{Opt}^{K-1}_{\{1,\ldots,k-1\}}$ using Algorithm 1
  Let $U = \max\{f_k(y) | y \in \mathrm{Opt}^{K-1}_{\{1,\ldots,k-1\}}\}$
  Let $L = \mathrm{Opt}^1_k$
  Let $step = \frac{U-L}{t}$
  for $t \in \{0, \ldots, t-1\}$ **do**
    Let $l = L + t \times step$
    Let $u = l + step$
    Start a MOIP solver in a new thread to find all solutions y satisfying $l < f_k(y) \leq u$.
  end
  Return the union of the results from all threads started
end

Algorithm 1: The Efficient Projection Partitioning (EPP) algorithm.

3.3. **K-PPM.** We now give a brief overview of the algorithm K-PPM from [7]. The algorithm is recursive, and determining Opt^K depends on knowing the results of Opt^{K-1}_S for all subsets $S \subset \{1, \ldots, K\}$ of size $K-1$. The solutions to these sub-problems are used to determine the ideal and nadir vectors for Opt^K. These are used in turn to divide the search space.

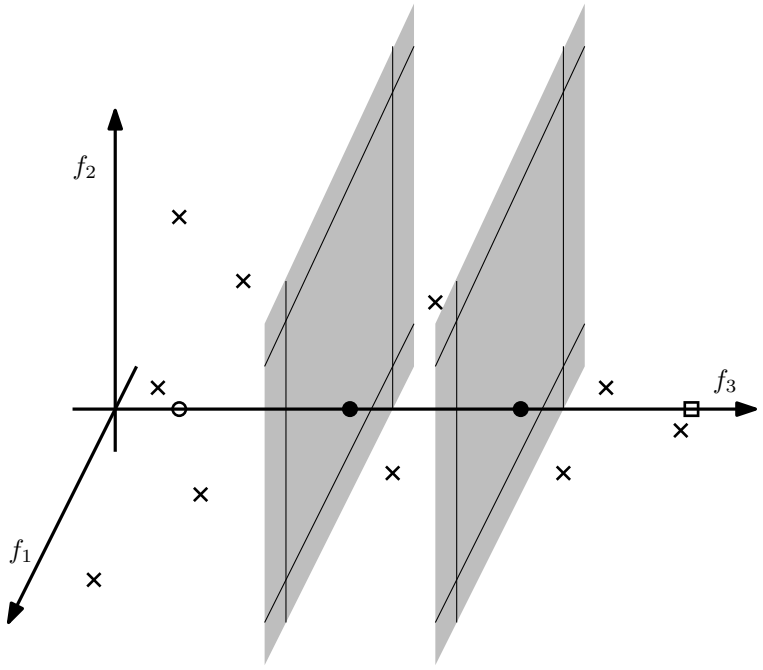

FIGURE 1. A visualisation of EPP splitting a 3-objective problem. The grey planes extend to infinity in the f_1 and f_2 directions. The hollow circle denotes the minimal value of f_3, and the hollow square marks the maximal value of f_3. The filled circles mark out equal divisions of this range, which are also the intersections of the grey planes with the f_3 axis.. Each of the crosses marks a solution to the MOIP. This partitioning will launch 3 worker threads, one for each partition.

For each objective bar one, the range of values that objective can take is divided into a number of blocks. By taking one such block from each objective, a "column" is formed that has upper and lower bounds on each dimension bar one. Then the algorithm solves the single objective optimisation problem for this final objective.

This gives a set of well-distributed solutions. These well-distributed solutions are then used to create "boxes" in which to search for all Pareto solutions. An initial box is created by taking the ideal and nadir vectors as two opposing corners. Then each well-distributed solution is used to partition the box which contains it. The new solution is used as a central point which, when combined with the local ideal and nadir vectors, forms 2^K new boxes. Of these new boxes, two are immediately discarded. One is discarded as the solution being inserted would dominate any possible solution in this region. The second is discarded as any solutions contained therein would dominate the given solution, a contradiction.

For each box, the algorithm creates a sub-problem with the addition of upper and lower bounds from the box. Then a MOIP solver finds all solutions to this sub-problem, and the set Opt^K is the union of solutions found by solvers running across all boxes.

K-PPM can create a rather large number of sub-problems, especially in higher dimensions. There is also a strict hierarchy that dictates that certain problems should be solved before others (i.e. Opt^2 must be known before Opt^3 can be started). As a result, K-PPM does not tie the number of threads to the partitioning process, but instead launches some number of threads, and uses a "master" node to control which jobs are passed to which threads.

3.4. **CPLEX.** Unless configured otherwise, CPLEX will launch up to one thread per core available on a given machine. As MOIP-AIRA uses CPLEX as a single-objective IP solver, we chose to run MOIP-AIRA with multiple threads to also compare the parallelisation of CPLEX. This is not a partitioning algorithm, it merely uses more threads to solve each individual single-objective problem. Note that for both EPP and K-PPM, CPLEX was configured to only use one thread per MOIP-AIRA instance.

4. Implementation

Both EPP and K-PPM were implemented in C++11, using the shared memory and threading of the Standard Template Library to handle all thread creation and data sharing. MOIP-AIRA was used to solve sub-problems, as the availability of the source code[13] greatly sped up the implementation process.

IBM ILOG CPLEX [15] 12.7 was used as the single-objective solver for all algorithms, and CPLEX was limited to using 1 thread in partitioning algorithms. The code for EPP was integrated back into [13] while the code for K-PPM was kept separately at [16]. Source code for both of these are available publicly for further research.

For testing, a number of knapsack, assignment and travelling salesman problems with 3 and 4 objectives and with varying sizes were generated. For each problem type and problem size, 5 instances were randomly generated. This follows the work in [10] where random knapsack problems were solved. For those wishing to compare algorithms, these test cases are also provided at [17, 18].

All tests were run on National Computational Infrastructure facilities, which utilise Intel Xeon E5-2670 CPUs running at 2.60GHz with hyper-threading disabled. All code was compiled with GCC 4.9.0 using optimisation level `-O2`. For each algorithm, each set of five 3-objective problems was run with 3 and 6 threads, while each set of five 4-objective problems was run with 4, 8 and 12 threads. The average of the running time of each set of five problem are reported in this paper. Additionally, MOIP-AIRA was used as a baseline solver for each problem, only utilising one thread. This was done to show how well each algorithm performed when directly trying to improve the running time through parallelisation. Note that timing results from all problems from our data-sets are not necessarily presented. We excluded problems where algorithms took longer than 48 hours to complete (and was therefore terminated), as well as problems which were solved by most algorithms in under 10 seconds.

The source code for all algorithms tested [13, 16], as well as all test cases [17, 18], are available for further use.

5. Discussion

All tables clearly show that EPP outperforms KPPM on a raw timing basis. We also see that both algorithms do benefit from introducing more threads.

For some smaller problems K-PPM using 4 threads is actually slower than not partitioning the search space at all. This is at the same time noteworthy, and also not completely unexpected. To partition a 4-objective problem, K-PPM must calculate the nadir vector of the problem. This involves determining Opt_S^{K-1} for all subsets $S \subset \{1, \ldots, K\}$ of size $K-1$. This recursive procedure continues down, with K-PPM requiring all $\binom{K}{K-1}$ $(K-1)$-objective problems to be completely solved before partitioning each K-objective problem. There is some repeated use of the same solution sets, but this is still extra work that is required. As mentioned in [7], the calculation of this nadir vector can often completely dominate the running time of K-PPM, and the removal of this calculation can improve the running time but at the cost of losing exactness in the algorithm.

In contrast, EPP only requires Opt_S^{K-1} for one subset S, and can still calculate the Pareto set exactly. In total, this means that on any given 4-objective problem, EPP will solve 3 fewer 3-objective problems, 5 fewer 2-objective problems and 3 fewer 1-objective problems. The timing results indicate that as more objectives are introduced, EPP will scale better than K-PPM.

A large number of small sub-problems can also adversely affect a MOIP solver. MOIP-AIRA achieves significant speed improvements when it can take advantage of earlier solved problems, as discussed in [12]. Since K-PPM divides the solution space along all-bar-one objectives in its first phase, it will find at least 2^{K-1} well-distributed solutions to partition the search space. Each such solution will take one box and divide it into $2^K - 2$ new boxes, resulting in at least $2^{2K-1} - 3 \cdot 2^{K-1}$ boxes to search. For $K = 4$, this is a minimum of 104 boxes. It is entirely likely that multiple small sub-problems will repeat the same calculations independently.

In contrast, EPP only creates one partition for each thread allocated to solve the problem. This does reduce the repetition of calculations. However, it also highlights an immediate drawback: there is no guarantee that each sub-problem will take exactly the same amount of time. Experimental results confirm that one thread will often take longer than all others, but EPP still performs faster than K-PPM.

5.1. Further work.

We believe there is potential for improving EPP even further by more accurately ensuring that all threads perform an equal amount of work. This could be achieved by establishing methods of correctly determining the specifics of the distribution of the objective values. An easier but less effective approach would be to investigate a relaxation of the relationship between number of threads and number of partitions. By using more partitions than threads, the work done by each thread could be more equally distributed, at a possible cost of repeated calculations.

The literature also currently only considers partitioning algorithms with independent workers. We question whether the performance of EPP (or indeed any other algorithm) could be improved by sharing the correct data between workers while they are still running. We see from our timing results that

	Average		K-PPM		CPLEX		EPP	
	number of	Single	3	6	3	6	3	6
Vars.	solutions	thread	threads	threads	threads	threads	threads	threads
10	201	18.61	15.66	9.53	11.09	11.01	9.39	6.37
15	616	100	72.56	48.28	61.82	58.74	47.94	29.52
20	1776	405	290	204	272	251	202	122
25	3398	1085	792	566	755	681	535	315
30	6571	2706	1818	1399	2743	1713	1396	801
40	12529	7051	4998	4056	6330	4624	4016	2183

TABLE 1. Running times of each algorithm and various thread counts on a number of assignment problems with 3 objective functions.

	Average		K-PPM		CPLEX		EPP	
	number of	Single	3	6	3	6	3	6
Objects	solutions	thread	threads	threads	threads	threads	threads	threads
50	342	38.62	34.98	21.69	38.64	38.90	21.78	13.62
75	1057	237	222	148	186	179	137	76.44
100	1855	667	572	454	491	461	412	259
125	4255	2063	1644	1189	1255	1102	1151	727
150	5358	3338	2524	2284	4223	1883	1816	1099
200	16951	18643	13176	11087	9118	8331	9946	6897

TABLE 2. Running times of each algorithm and various thread counts on a number of knapsack problems with 3 objective functions.

	Average		K-PPM		CPLEX		EPP	
	number of	Single	3	6	3	6	3	6
Cities	solutions	thread	threads	threads	threads	threads	threads	threads
12	126	14.21	13.65	9.31	12.41	10.45	8.85	5.29
15	734	164	129	88.78	143	111	93.99	58.63
20	1988	1026	641	467	748	752	569	323
30	10697	12213	8242	6058	9413	8451	6340	3529

TABLE 3. Running times of each algorithm and various thread counts on a number of travelling salesman problems with 3 objective functions.

	Average		K-PPM			CPLEX			EPP		
	number of	Single	4	8	12	4	8	12	4	8	12
Vars.	solutions	thread	thr.	thr.	thr.	thr.	thr.	thr.	thr.	thr.	thr.
8	264	67.66	121	61.52	45.35	37.72	36.60	1097	31.72	22.12	16.11
10	1098	532	732	328	241	242	230	4493	240	125	97.55
11	1917	1095	1411	642	471	513	476	7858	539	279	201
12	2104	1284	1567	732	533	621	571	5972	634	325	246
15	8689	7018	7149	3655	2757	3742	3061	21736	3718	1779	1287
20	40298	51505	29137	17060	14499	31234	22761	-	24788	12834	9126

TABLE 4. Running times of each algorithm and various thread counts on a number of assignment problems with 4 objective functions.

	Average		K-PPM			CPLEX			EPP		
	number of	Single	4	8	12	4	8	12	4	8	12
Objects	solutions	thread	thr.	thr.	thr.	thr.	thr.	thr.	thr.	thr.	thr.
40	590	177	275	148	111	169	153	1932	116	63.24	56.16
60	3903	2529	2329	1205	1046	2223	2242	26434	1564	776	600
80	9379	17265	15095	8326	6617	6774	8429	-	10082	7251	4089

TABLE 5. Running times of each algorithm and various thread counts on a number of knapsack problems with 4 objective functions.

	Average		K-PPM			CPLEX			EPP		
	number of	Single	4	8	12	4	8	12	4	8	12
Cities	solutions	thread	thr.	thr.	thr.	thr.	thr.	thr.	thr.	thr.	thr.
8	96	19.58	48.50	27.39	19.87	14.67	16.09	1830	9.67	6.32	5.14
10	530	301	391	221	159	212	190	3369	148	82.28	64.99
12	1541	1502	1662	990	751	1138	928	8918	819	470	345
15	5613	12982	10558	6011	5473	9255	7754	-	7024	4121	3036

TABLE 6. Running times of each algorithm and various thread counts on a number of travelling salesman problems with 4 objective functions.

no algorithm achieves perfect scaling as more threads are used, so improvement is definitely possible. Sharing data "online" while threads are still running has the potential to reduce the number of redundant calculations, but does come at increased complexity.

6. Conclusion

This paper introduced a new method of partitioning the search space of a MOIP, called the Efficient Projection Partitioning (EPP) algorithm. By projecting solutions onto a single dimension and only dividing the solution space along this one dimension, EPP avoids the expensive calculation of the nadir vector which dominates the running time of other state-of-the-art algorithms. EPP is shown to be faster on every single test case examined, including being able to scale well on low numbers of threads when compared to non-parallel algorithms. The performance and scaling of EPP with low numbers of threads makes it much more feasible to solve smaller problems on consumer-friendly devices such as desktops, laptops or tablets. As EPP is also faster on larger problems, and when more threads are introduced, EPP will allow users to exactly solve multi-objective integer programming problems with more variables or more objectives. The source code to our algorithm is provided to help the development and utilisation of this work.

7. Acknowledgements

This study was supported by the Australian Research Council under the Discovery Projects funding scheme (project DP140104246).

References

[1] K. E. Parsopoulos, M. N. Vrahatis, Particle swarm optimization method in multiobjective problems, in: Proceedings of the 2002 ACM Symposium on Applied Computing, SAC '02, ACM, New York, NY, USA, 2002, pp. 603–607. doi:10.1145/508791.508907.

[2] K. Sethanan, W. Neungmatcha, Multi-objective particle swarm optimization for mechanical harvester route planning of sugarcane field operations, European Journal of Operational Research 252 (3) (2016) 969 – 984. doi:http://dx.doi.org/10.1016/j.ejor.2016.01.043.
URL http://www.sciencedirect.com/science/article/pii/S0377221716000886

[3] E. Cantu-Paz, Efficient and accurate parallel genetic algorithms, Vol. 1, Springer Science & Business Media, 2000.

[4] P. M. Reed, J. B. Kollat, M. P. Ferringer, T. G. Thompson, Parallel evolutionary multi-objective optimization on large, heterogeneous clusters: An applications perspective, Journal of Aerospace Computing, Information and Communications 5 (11) (2008) 460–478. doi:10.2514/1.36467.

[5] M. Ehrgott, X. Gandibleux, A survey and annotated bibliography of multiobjective combinatorial optimization, OR-Spektrum 22 (4) (2000) 425–460. doi:10.1007/s002910000046.
URL http://dx.doi.org/10.1007/s002910000046

[6] A. Przybylski, X. Gandibleux, Multi-objective branch and bound, European Journal of Operational Research (2017) –doi:http://dx.doi.org/10.1016/j.ejor.2017.01.032.
URL http://www.sciencedirect.com/science/article/pii/S037722171730067X

[7] C. Dhaenens, J. Lemesre, E. Talbi, K-ppm: A new exact method to solve multi-objective combinatorial optimization problems, European Journal of Operational Research 200 (1) (2010) 45 – 53. doi:http://dx.doi.org/10.1016/j.ejor.2008.12.034.
URL http://www.sciencedirect.com/science/article/pii/S0377221708010709

[8] J. Lemesre, C. Dhaenens, E. Talbi, Parallel partitioning method (ppm): A new exact method to solve bi-objective problems, Computers & Operations Research 34 (8) (2007) 2450 – 2462. doi:http://dx.doi.org/10.1016/j.cor.2005.09.014.
URL http://www.sciencedirect.com/science/article/pii/S0305054805003059

[9] N. Boland, H. Charkhgard, M. Savelsbergh, The triangle splitting method for biobjective mixed integer programming, in: International Conference on Integer Programming and Combinatorial Optimization, Springer, 2014, pp. 162–173.

[10] G. Kirlik, S. Sayın, A new algorithm for generating all nondominated solutions of multiobjective discrete optimization problems, European Journal of Operational Research 232 (3) (2014) 479 – 488. doi:http://dx.doi.org/10.1016/j.

ejor.2013.08.001.
URL http://www.sciencedirect.com/science/article/pii/S0377221713006474
[11] M. Laumanns, L. Thiele, E. Zitzler, An efficient, adaptive parameter variation scheme for metaheuristics based on the epsilon-constraint method, European Journal of Operational Research 169 (3) (2006) 932 – 942. doi:http://dx.doi.org/10.1016/j.ejor.2004.08.029.
URL http://www.sciencedirect.com/science/article/pii/S0377221704005715
[12] M. Ozlen, B. A. Burton, C. A. G. MacRae, Multi-objective integer programming: An improved recursive algorithm, Journal of Optimization Theory and Applications 160 (2) (2014) 470–482. doi:10.1007/s10957-013-0364-y.
[13] W. Pettersson, C. A. G. MacRae, moip_aira (2014-2017).
URL https://github.com/WPettersson/moip_aira
[14] M. Özlen, M. Azizoğlu, Multi-objective integer programming: A general approach for generating all non-dominated solutions, European Journal of Operational Research 199 (1) (2009) 25 – 35. doi:http://dx.doi.org/10.1016/j.ejor.2008.10.023.
URL http://www.sciencedirect.com/science/article/pii/S0377221708009624
[15] IBM, IBM ILOG CPLEX Optimization Studio (1988-2017).
URL https://cplex.com
[16] W. Pettersson, kppm (2017).
URL https://github.com/WPettersson/kppm
[17] W. Pettersson, 3 objective problem test casesdoi:10.6084/m9.figshare.4814695.v1.
URL https://figshare.com/articles/3_objective_problem_test_cases/4814695
[18] W. Pettersson, 4 objective problem test casesdoi:10.6084/m9.figshare.4814698.v1.
URL https://figshare.com/articles/4_objective_problem_test_cases/4814698

School of Science, RMIT University, Victoria, 3000, AUSTRALIA
E-mail address: william@ewpettersson.se

School of Science, RMIT University, Victoria, 3000, AUSTRALIA
E-mail address: melih.ozlen@rmit.edu.au